Redefining Performance Through System Balance

05/16/06

Spring Processor Forum

2006

Chuck Moore
AMD Senior Fellow

Talk Outline

- What drives our industry?
- Elements of Balanced System Performance
  - System-level bandwidth and power consumption
  - Single thread versus multi-thread performance
  - Understanding compatibility and transitions
  - Reliability
- AMD's Next Generation Processor Technology
  - Direct Connect Architecture enhancements
  - Core enhancements
- Crystal Ball: *New trends that may shape the future*
- Summary

Moore's Law (*Gordon Moore, Electronics Magazine – April 1965*)

What he said:

- Double transistor integration every 18 months

Another trend often confused with Moore's Law:

- Double *performance* every 18 months
- This is actually an example of ongoing *customer value*

The semiconductor industry is dependent upon ongoing customer value:

A virtuous cycle:

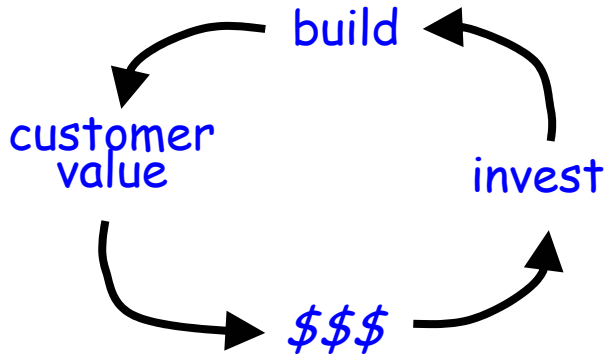

The Question we should all be asking ourselves:

How do we continue increasing customer value?

Customer Value Trends

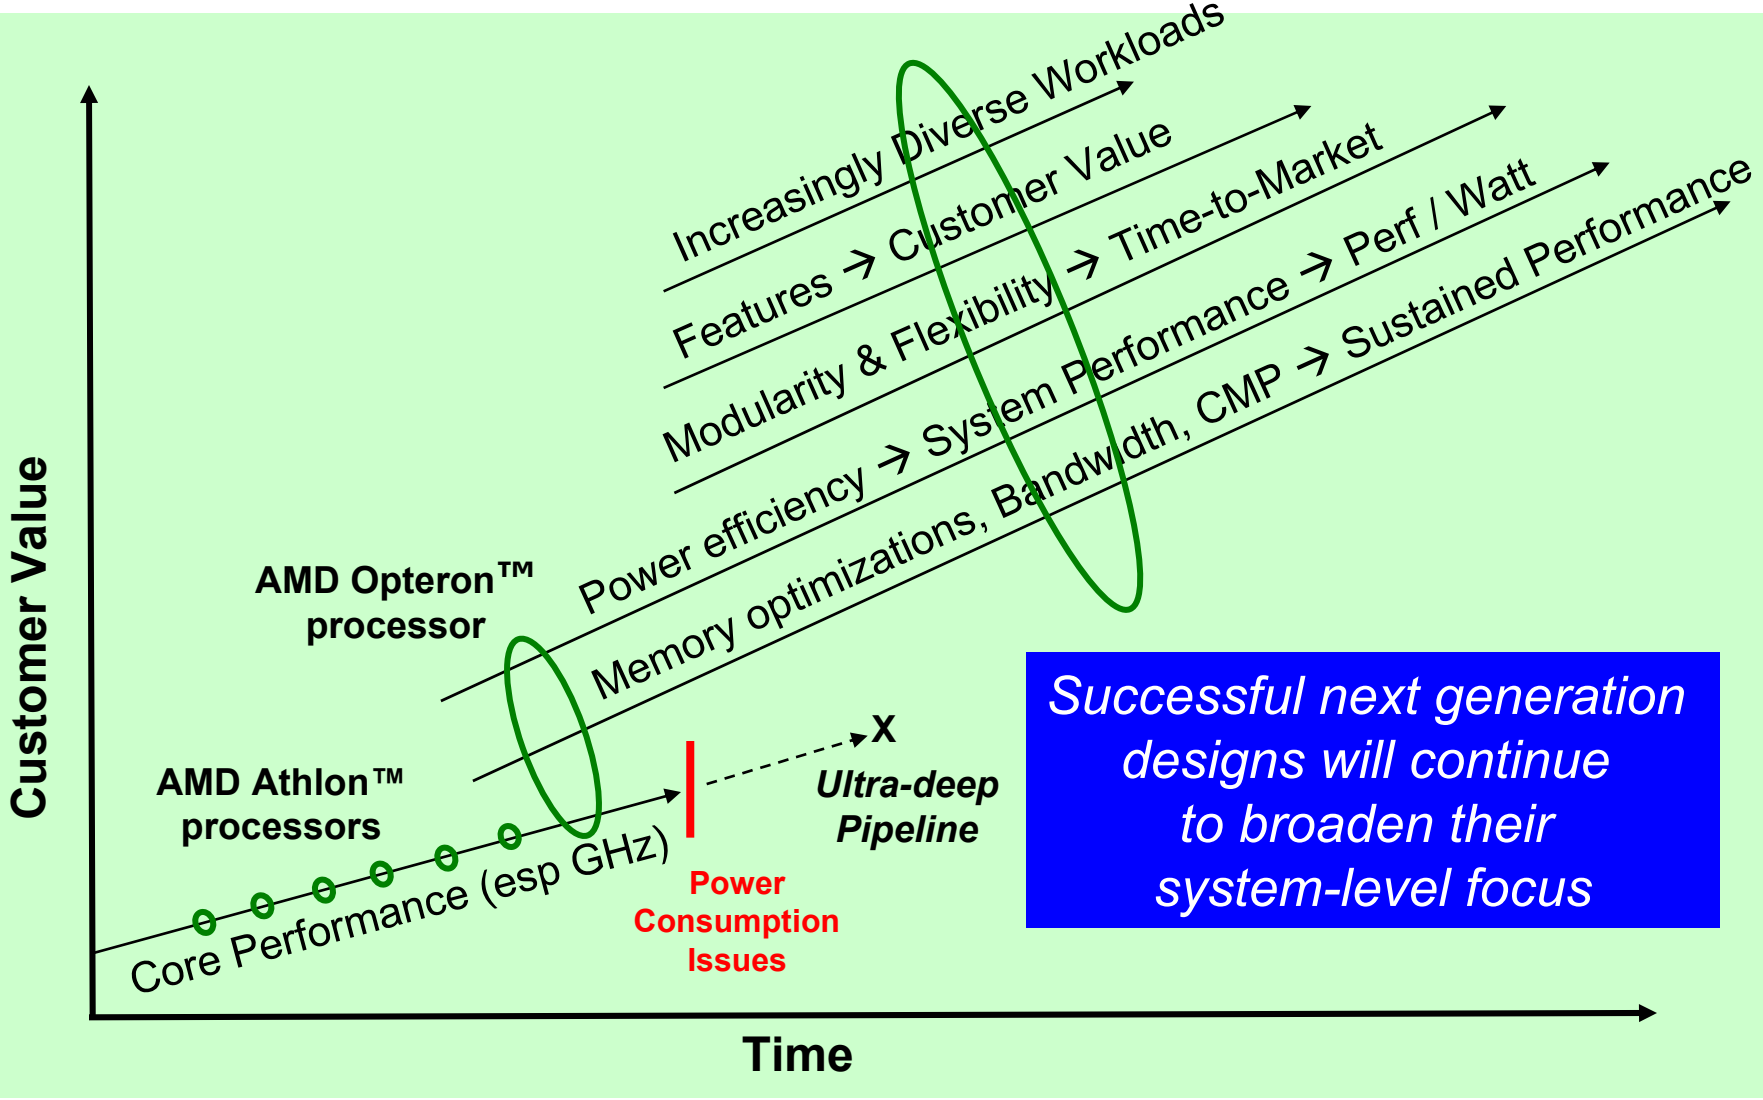

AMD is committed to understanding what customers *really* want

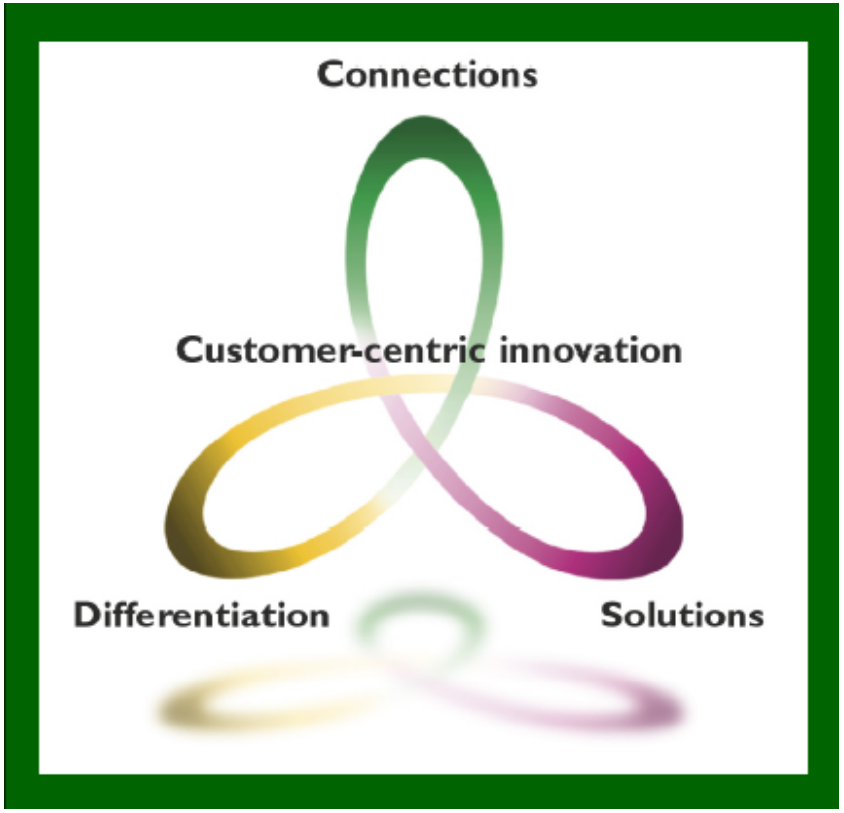

Expanding *connections* with customers, partners and end-users

Evolving beyond technologies and products to *solutions*

Enabling meaningful customer *differentiation* at ever declining costs

Customer-centric innovation

MPF 2004 - AMD Dual-Core Processor Chip

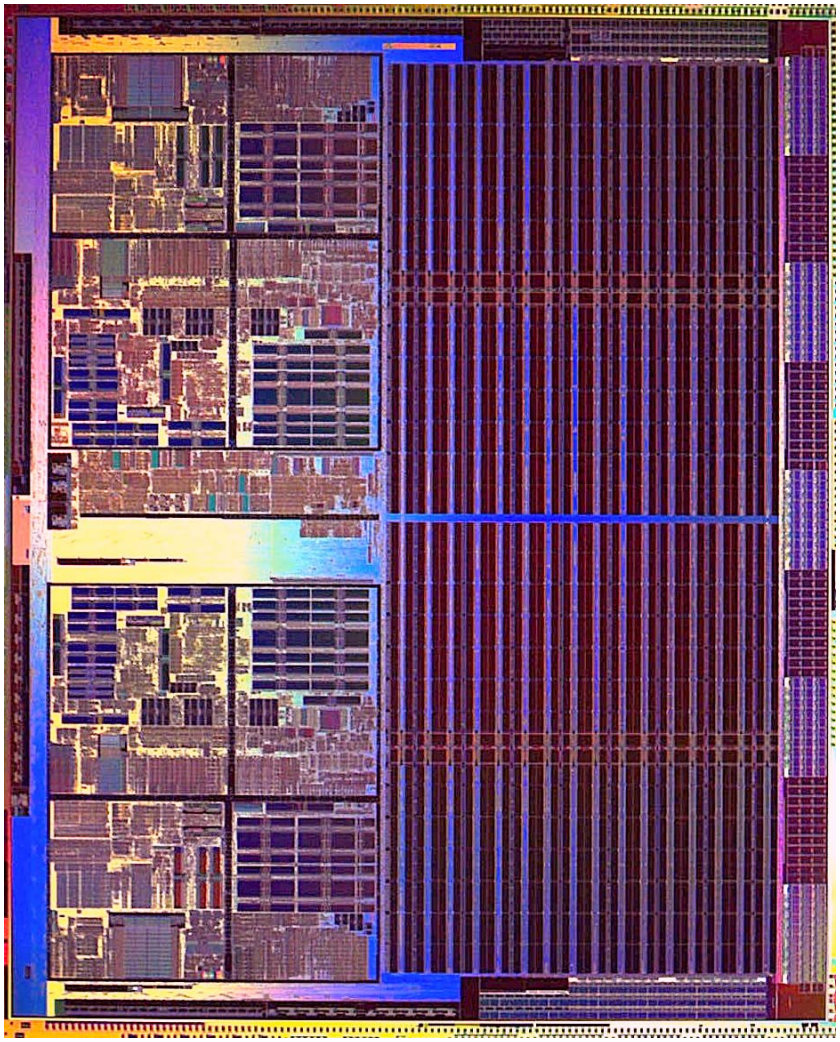

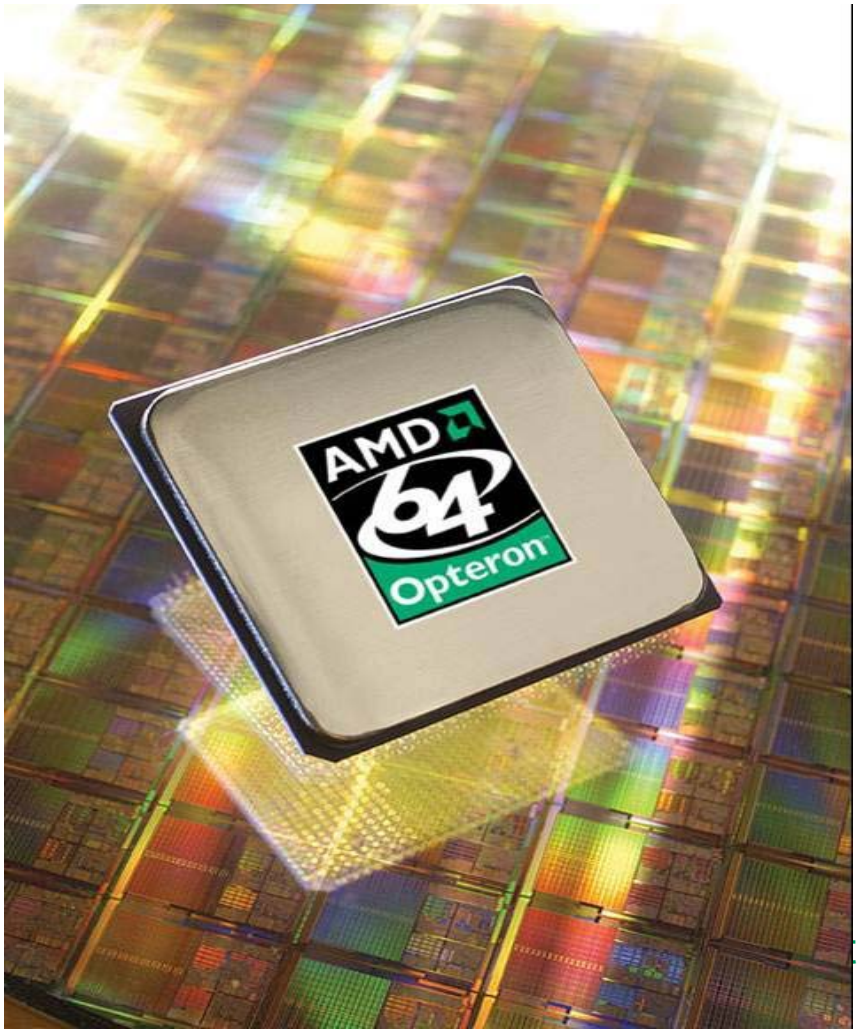

Talk Outline

- What drives our industry?
- Elements of Balanced System Performance
  - System-level bandwidth and power consumption
  - Single thread versus multi-thread performance
  - Understanding compatibility and transitions
  - Reliability
- AMD's Next Generation Processor Technology
  - Direct Connect Architecture enhancements
  - Core enhancements
- Crystal Ball: *New trends that may shape the future*
- Summary

Balanced Platform Bandwidth

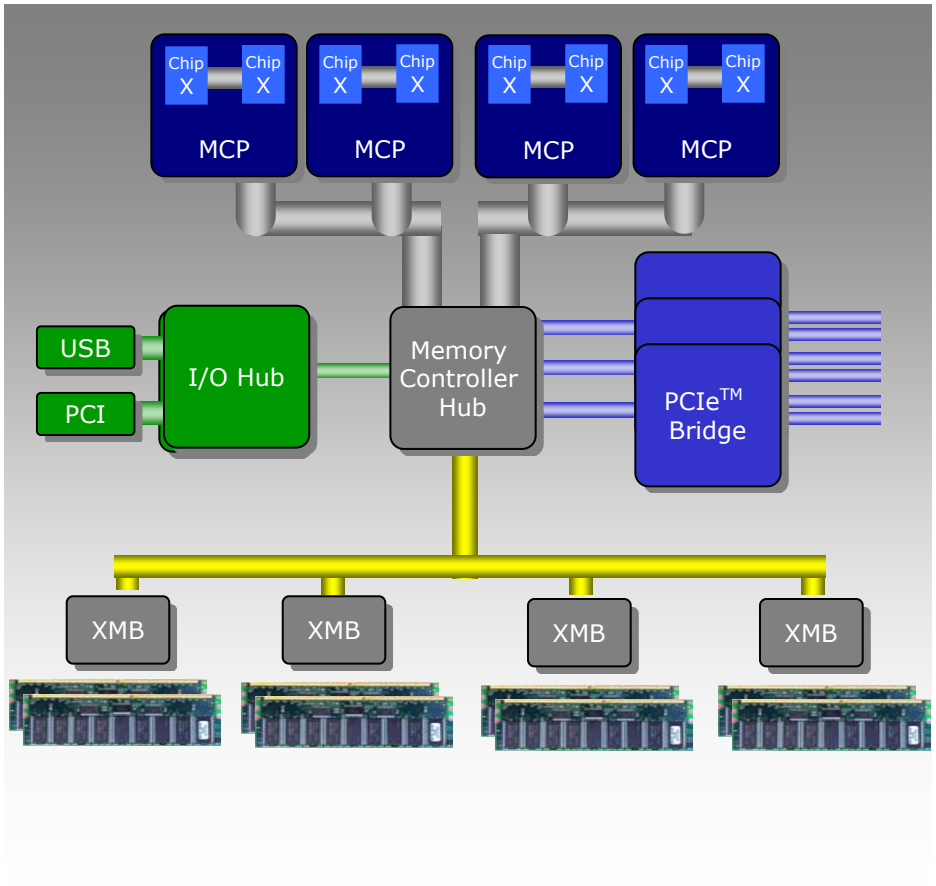

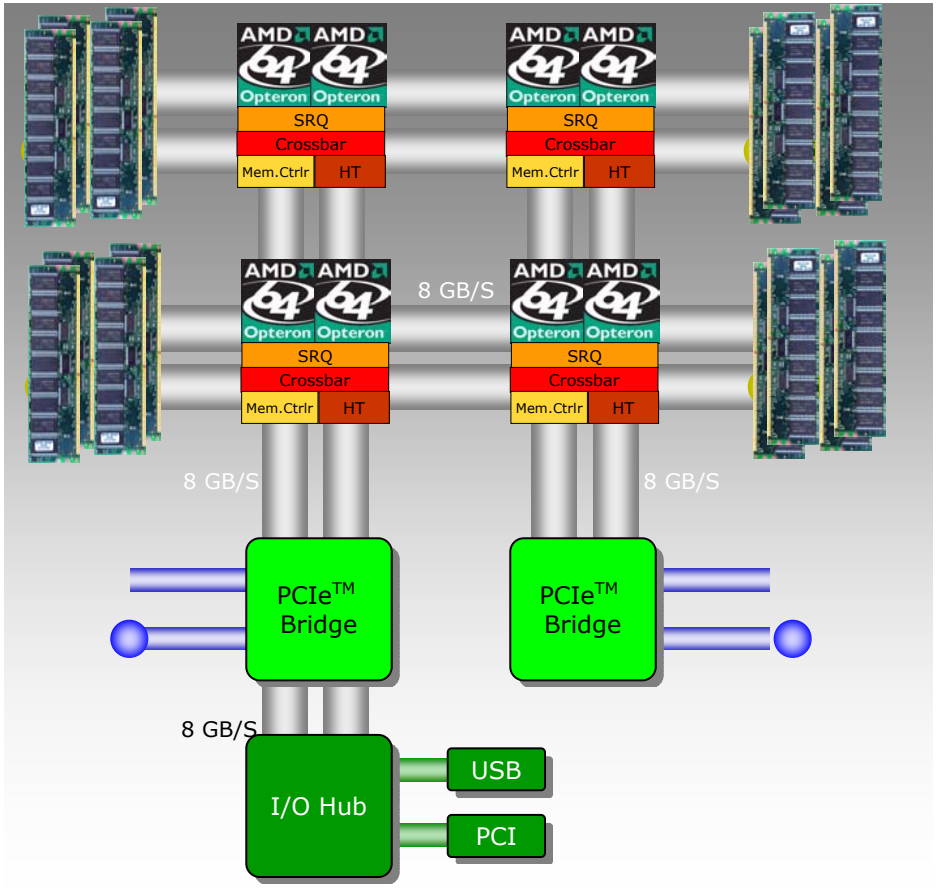

Legacy x86 Architecture

- 20-year old front-side bus (FSB) architecture
- CPUs, Memory, I/O all share a bus
- Major bottleneck to performance
- Faster CPUs or more cores ≠ performance

AMD64's Direct Connect Architecture

- Industry-standard technology
- Direct Connect helps eliminate the FSB bottleneck
- HyperTransport™ interconnect offers scalable high bandwidth and low latency

System-level Power Consumption – *Present Day*

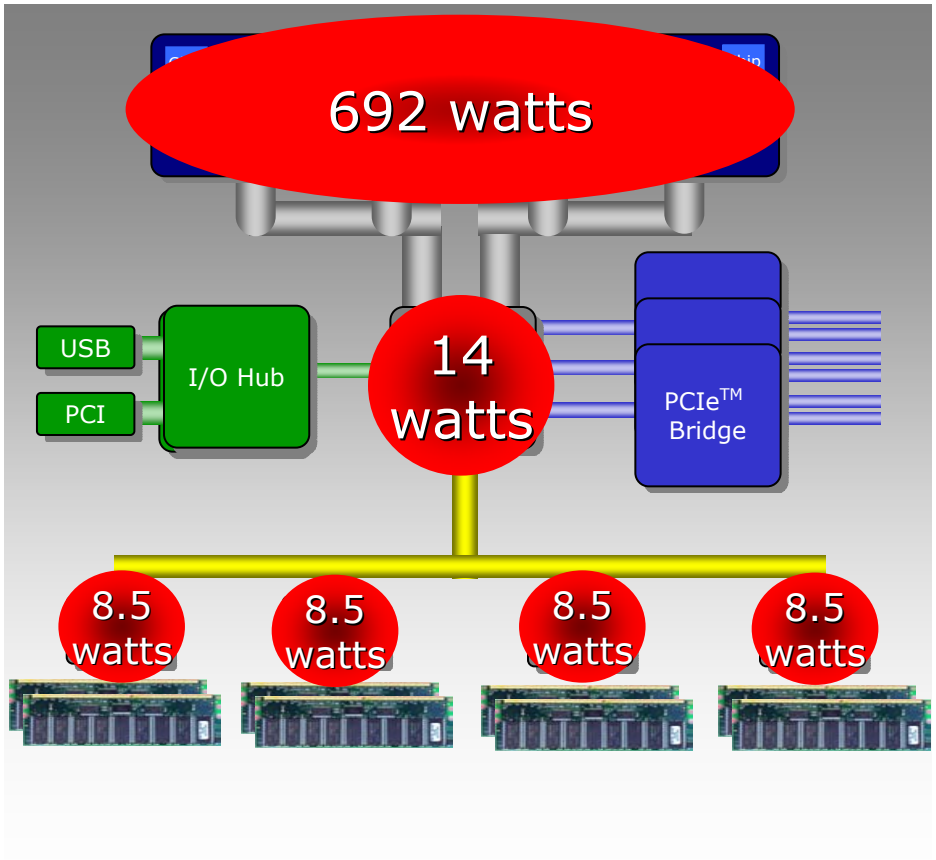

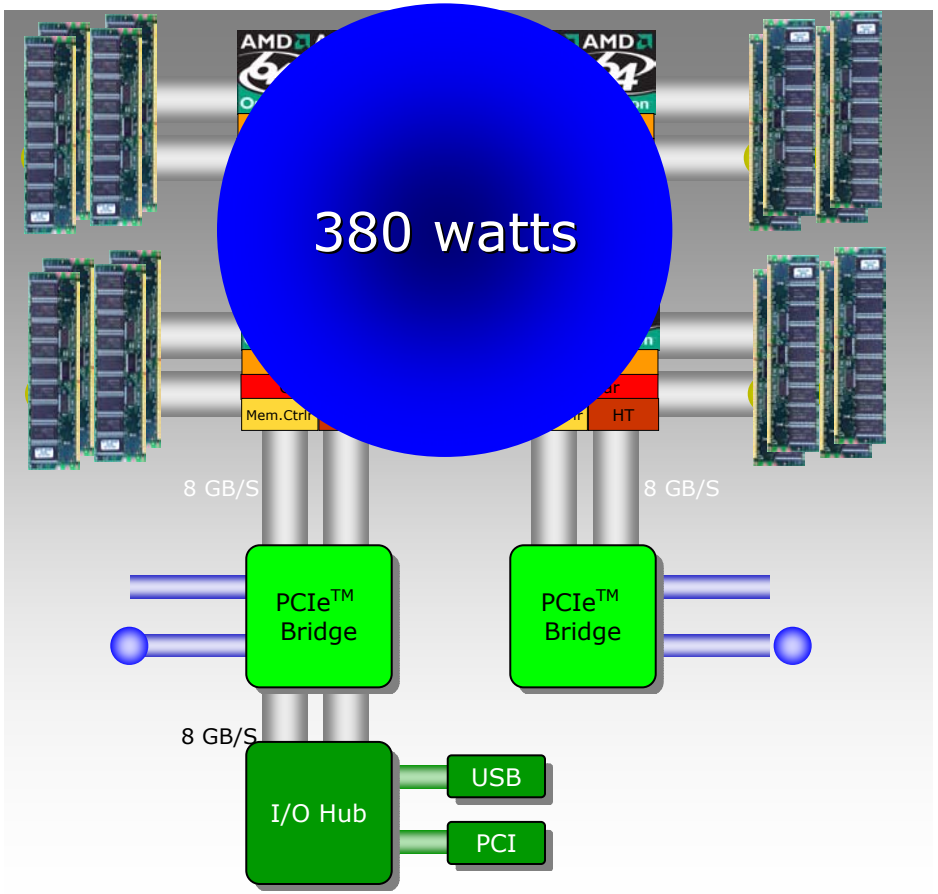

Dual-Core Packages with legacy technology

- 692 watts for processors (173w each)
- 48 watts for external memory controller

95% More Power

740 watts

Dual-Core AMD Opteron™ processors

- 380 watts for processors (95w each)
- Integrated memory controllers

380 watts

Source: Mixture of publicly available data sheets and AMD internal estimates.
Actual system power measurements may vary based on configuration and components used

AMD

Taking a look at FBDIMM - Effects on Power

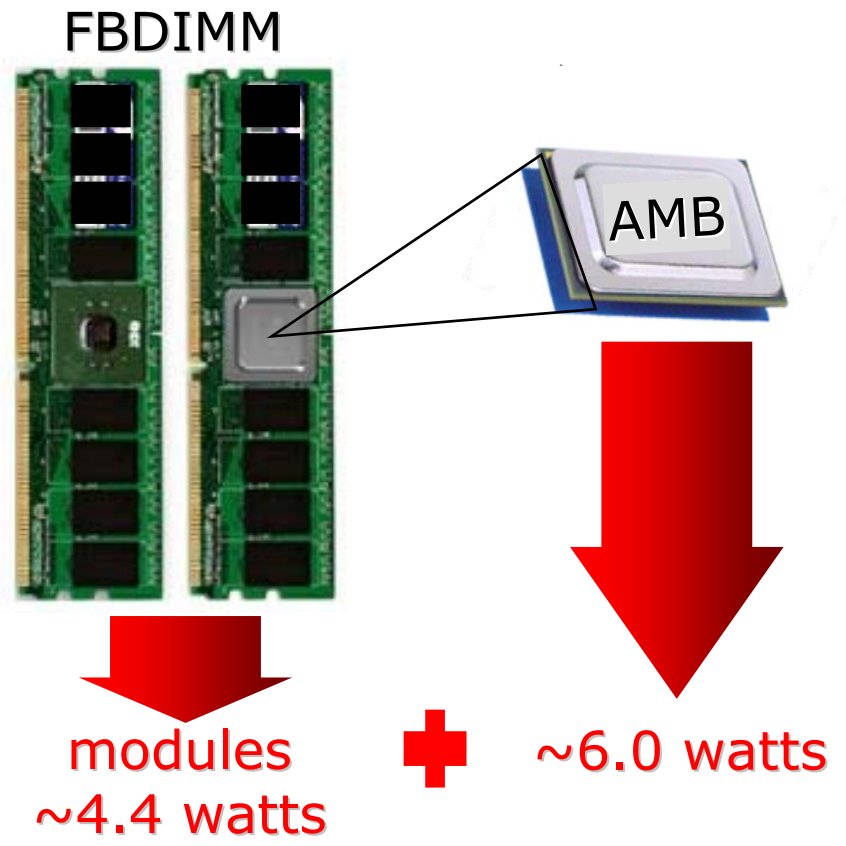

Pro's
- Increased Memory Capacity

Cons
- More Latency
- Increased Power

Example:

32 FBDIMM modules = ~333 watts

32 DDR2 modules = ~140 watts

FBDIMM = ~10.4 watts each

System-level Power Consumption (Plan) – *Soon*

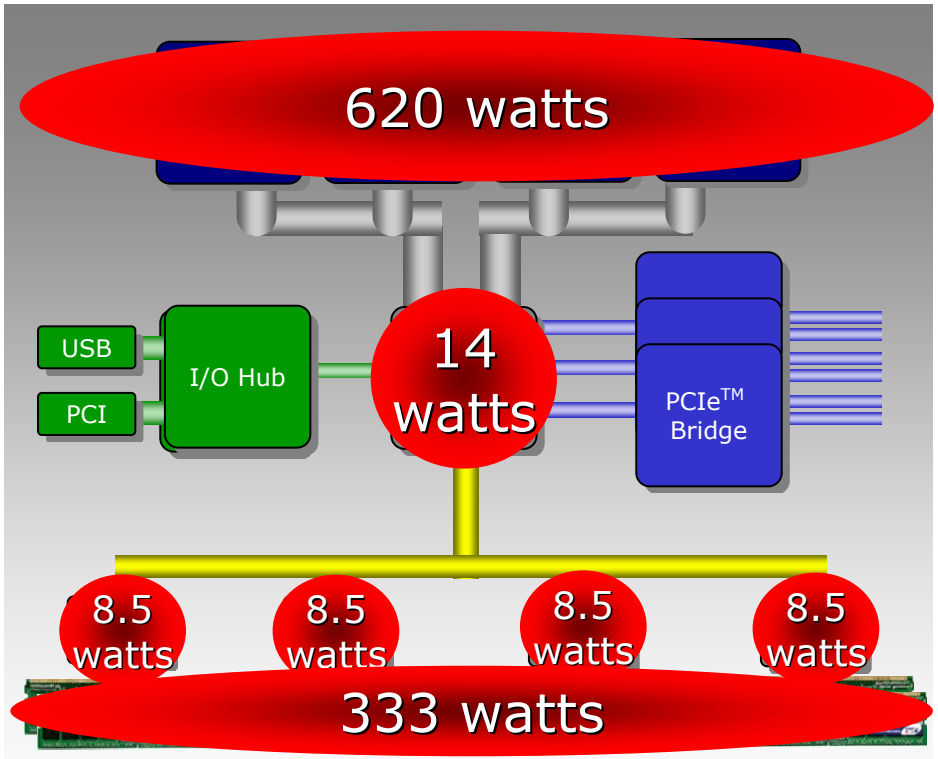

Legacy Dual-Core Chips in FBD Platform

- 620 watts for processors (155w each)
- 48 watts for external memory controller
- 333 watts for FBDIMM

93% More Power

1001 watts

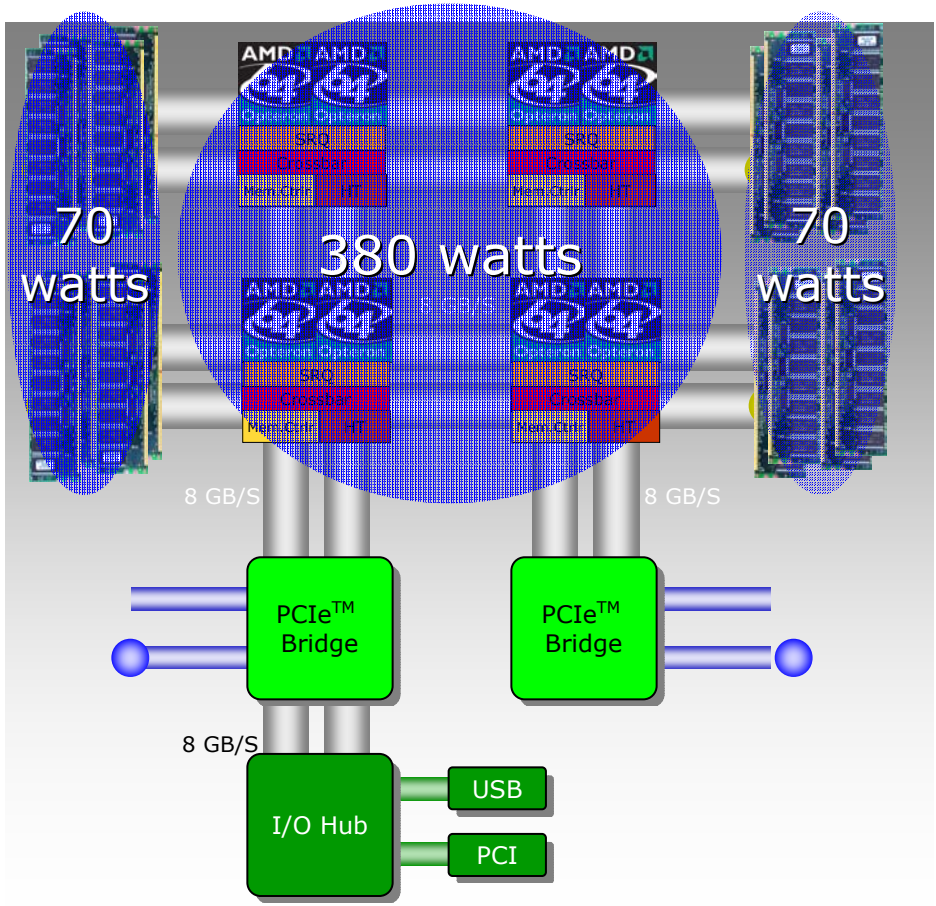

Dual-Core AMD Opteron™ processors

- 380 watts for processors (95w each)
- Integrated memory controllers
- 140 watts for DDR2

520 watts

Source: Mixture of publicly available data sheets and AMD internal estimates.
Actual system power measurements may vary based on configuration and components used

System-level Power Consumption (Plan) – *Later in 2006*

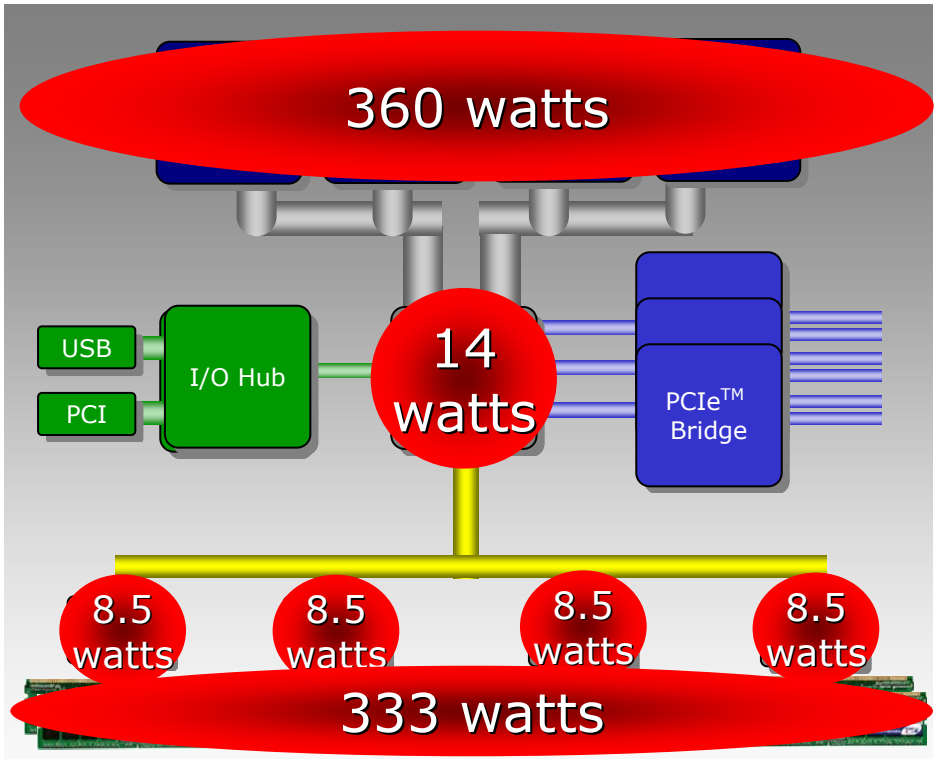

Improved Dual-Core Chips in FBD Platform

- 360 watts for processors (90w each)
- 48 watts for external memory controller
- 333 watts for FBDIMM

43% More Power

741 watts

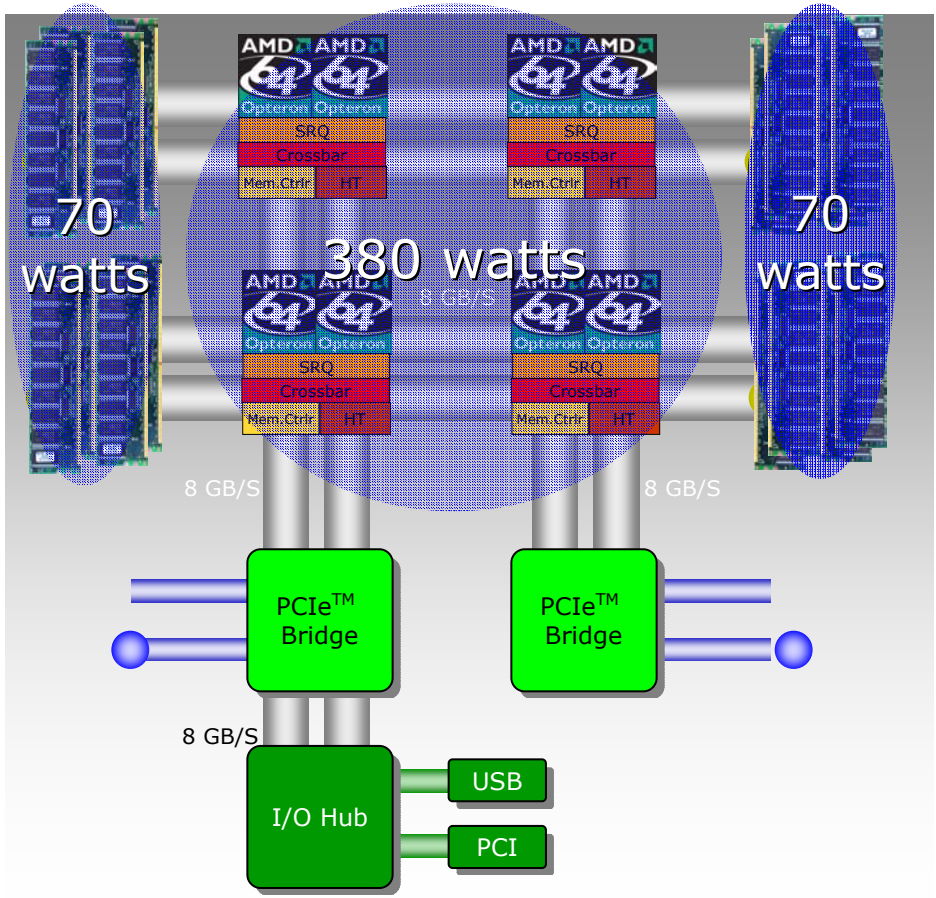

Dual-Core AMD Opteron™ processors

- 380 watts for processors (95w each)
- Integrated memory controllers
- 140 watts for DDR2

520 watts

Source: Mixture of publicly available data and AMD internal estimates.
Actual system power measurements may vary based on configuration and components used

AMD

What's Consuming all the Power?

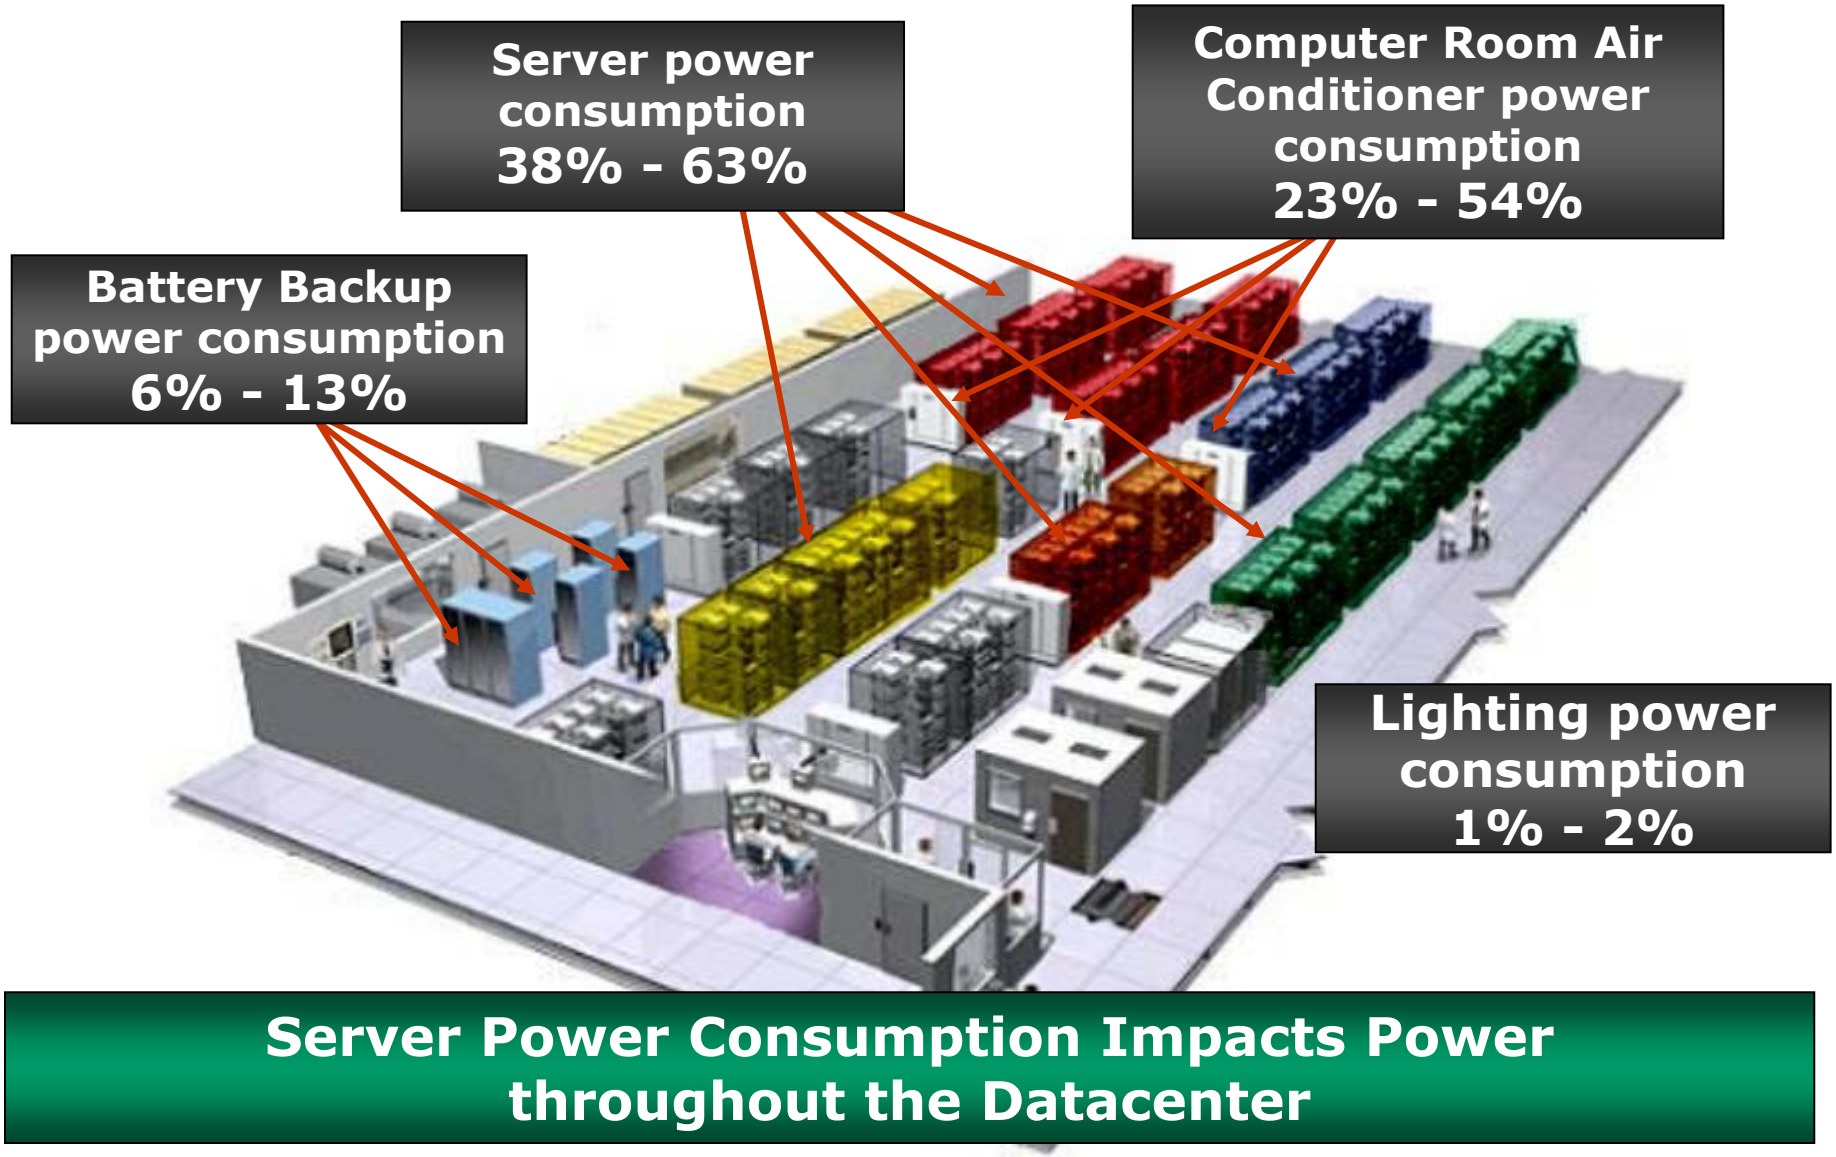

Server Power Consumption Impacts Power throughout the Datacenter

Reducing Power and Cooling Requirements with Processor Performance States

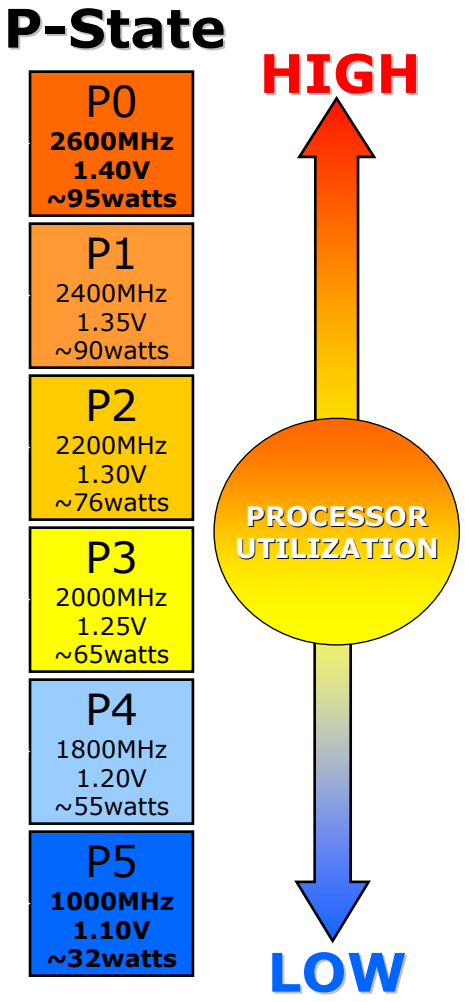

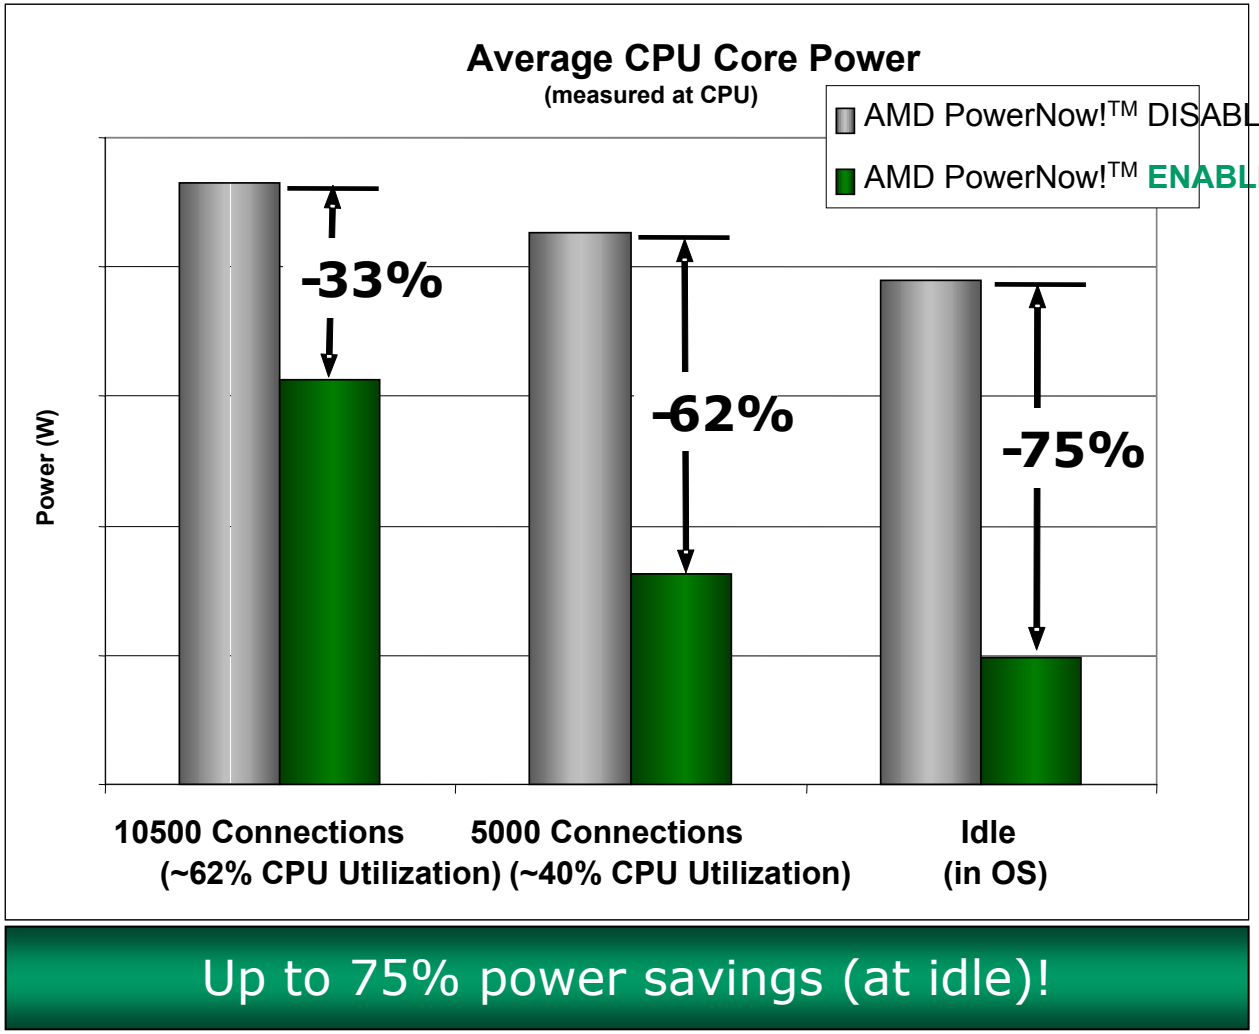

Up to 75% power savings (at idle)!

Talk Outline

- What drives our industry?
- Elements of Balanced System Performance
  - System-level bandwidth and power consumption
  - **Single thread versus multi-thread performance**
  - **Understanding compatibility and transitions**
  - **Reliability**
- AMD's Next Generation Processor Technology
  - Direct Connect Architecture enhancements
  - Core enhancements
- Crystal Ball: *New trends that may shape the future*
- Summary

A Closer Look at *Parallelism*

Instruction-level Parallelism (ILP)

- Executing multiple instructions from same program at the same time
- Superscalar hardware picks up most available ILP (*complexity effective*)

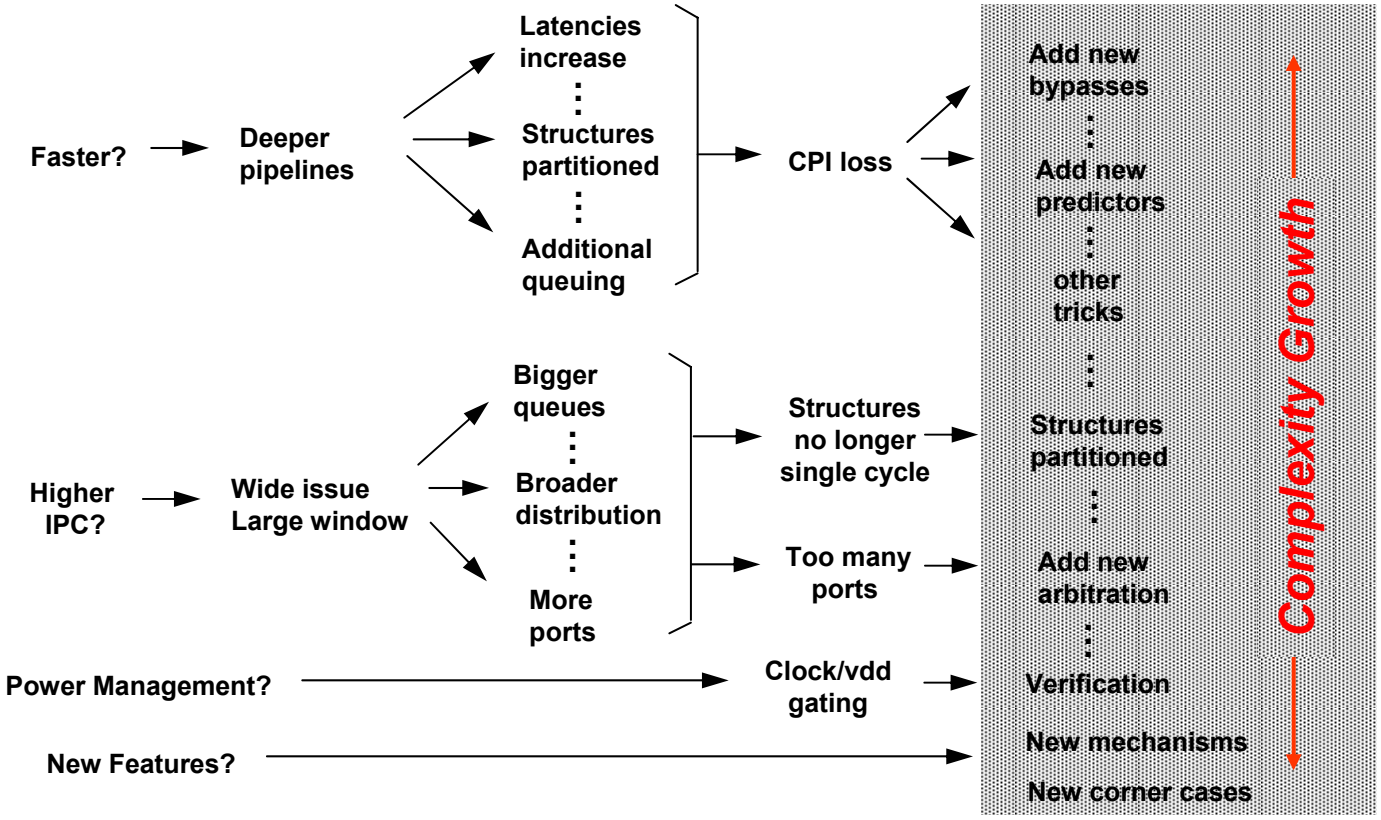

Superposition … *complexity tends to compound!*

A Closer Look at *Parallelism*

Instruction-level Parallelism (ILP)

- Executing multiple instructions from same program at the same time
- Superscalar hardware picks up most available ILP (*complexity effective*)

Data-level Parallelism (DLP)

- Executing same instruction on multiple pieces of data at the same time
- Vector-style processing -- SSE hardware operates in this manner

ADDSUBPS **Add and Subtract Packed Single-Precision**

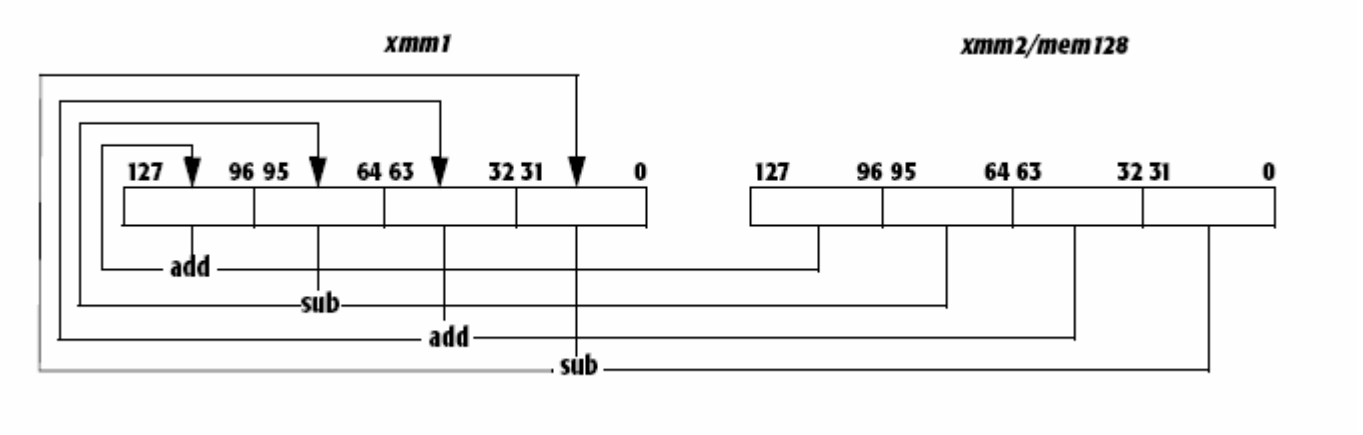

A Closer Look at *Parallelism*

Instruction-level Parallelism (ILP)

- Executing multiple instructions from same program at the same time
- Superscalar hardware picks up most available ILP (*complexity effective*)

Data-level Parallelism (DLP)

- Executing same instruction on multiple pieces of data at the same time
- Vector-style processing -- SSE hardware operates in this manner

Thread-level Parallelism (TLP) – *several types:*

1. Concurrent Applications
   - Multiple programs running at the same time
   - Multiple OS's on virtualized hardware image
   - Collection of services integrated into a single "application"
2. Internet Transactional
   - Multiple computers running the same application
3. Parallel Applications - *The holy grail of computer science*
   - Single application partitioned to run as multiple threads

CMP Performance *(Theoretical values)*

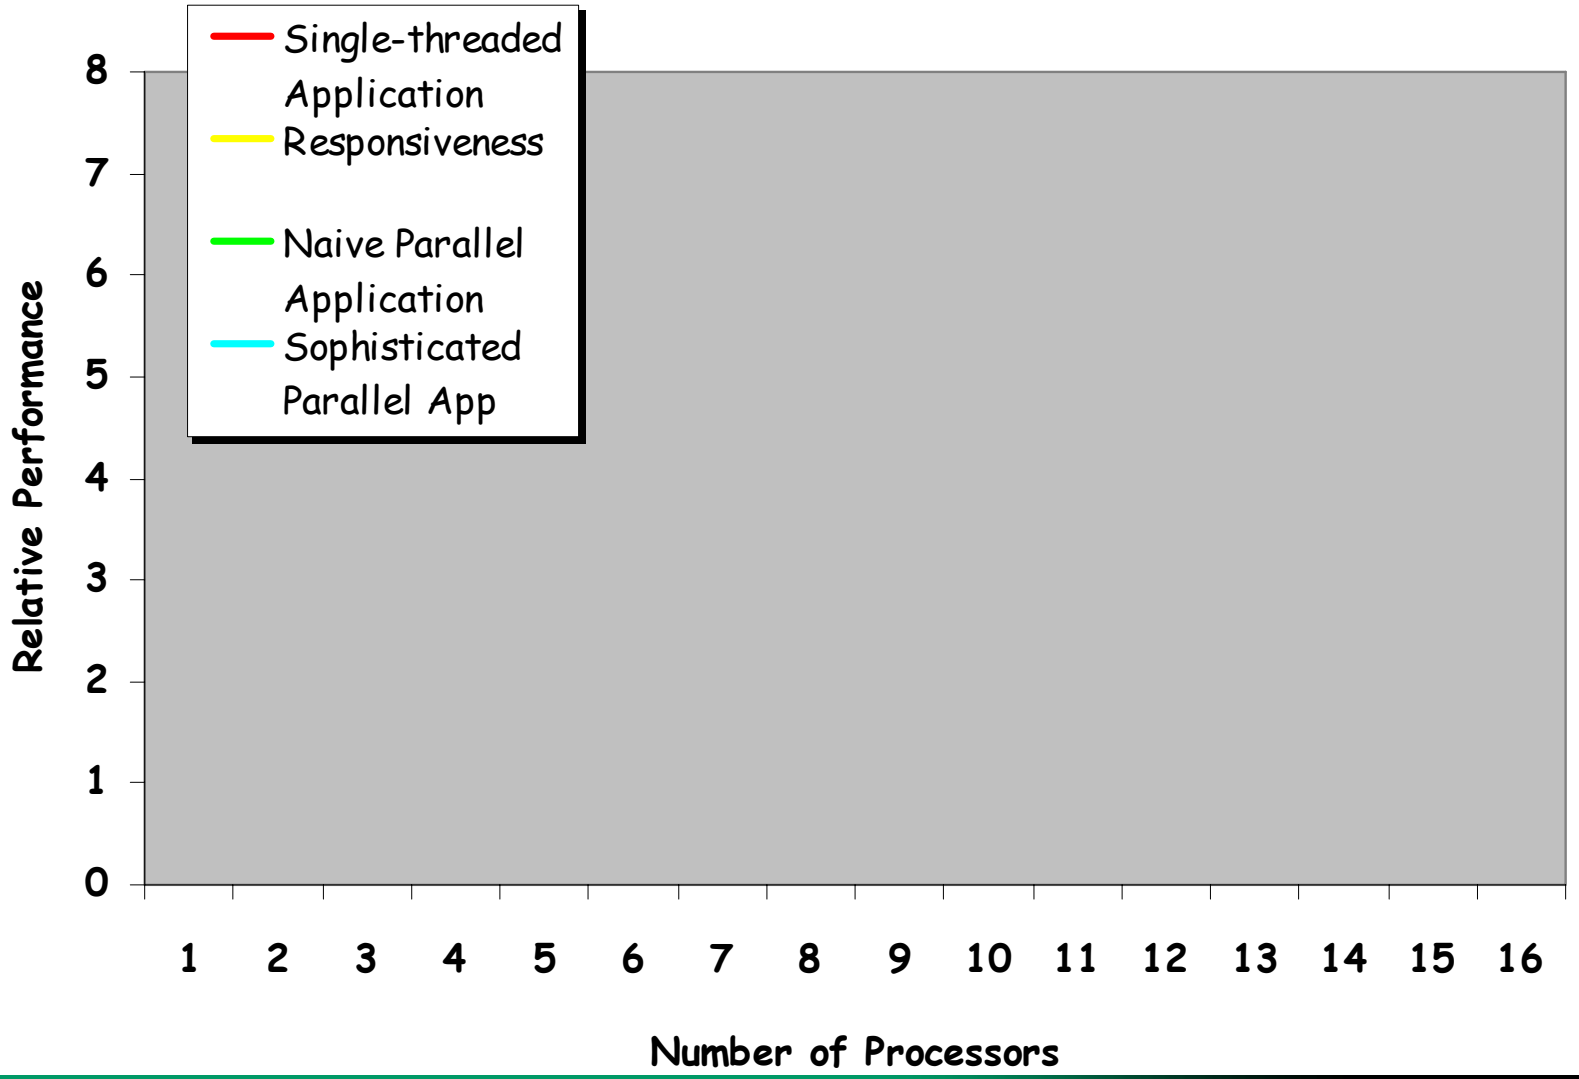

CMP Performance *(Theoretical values)*

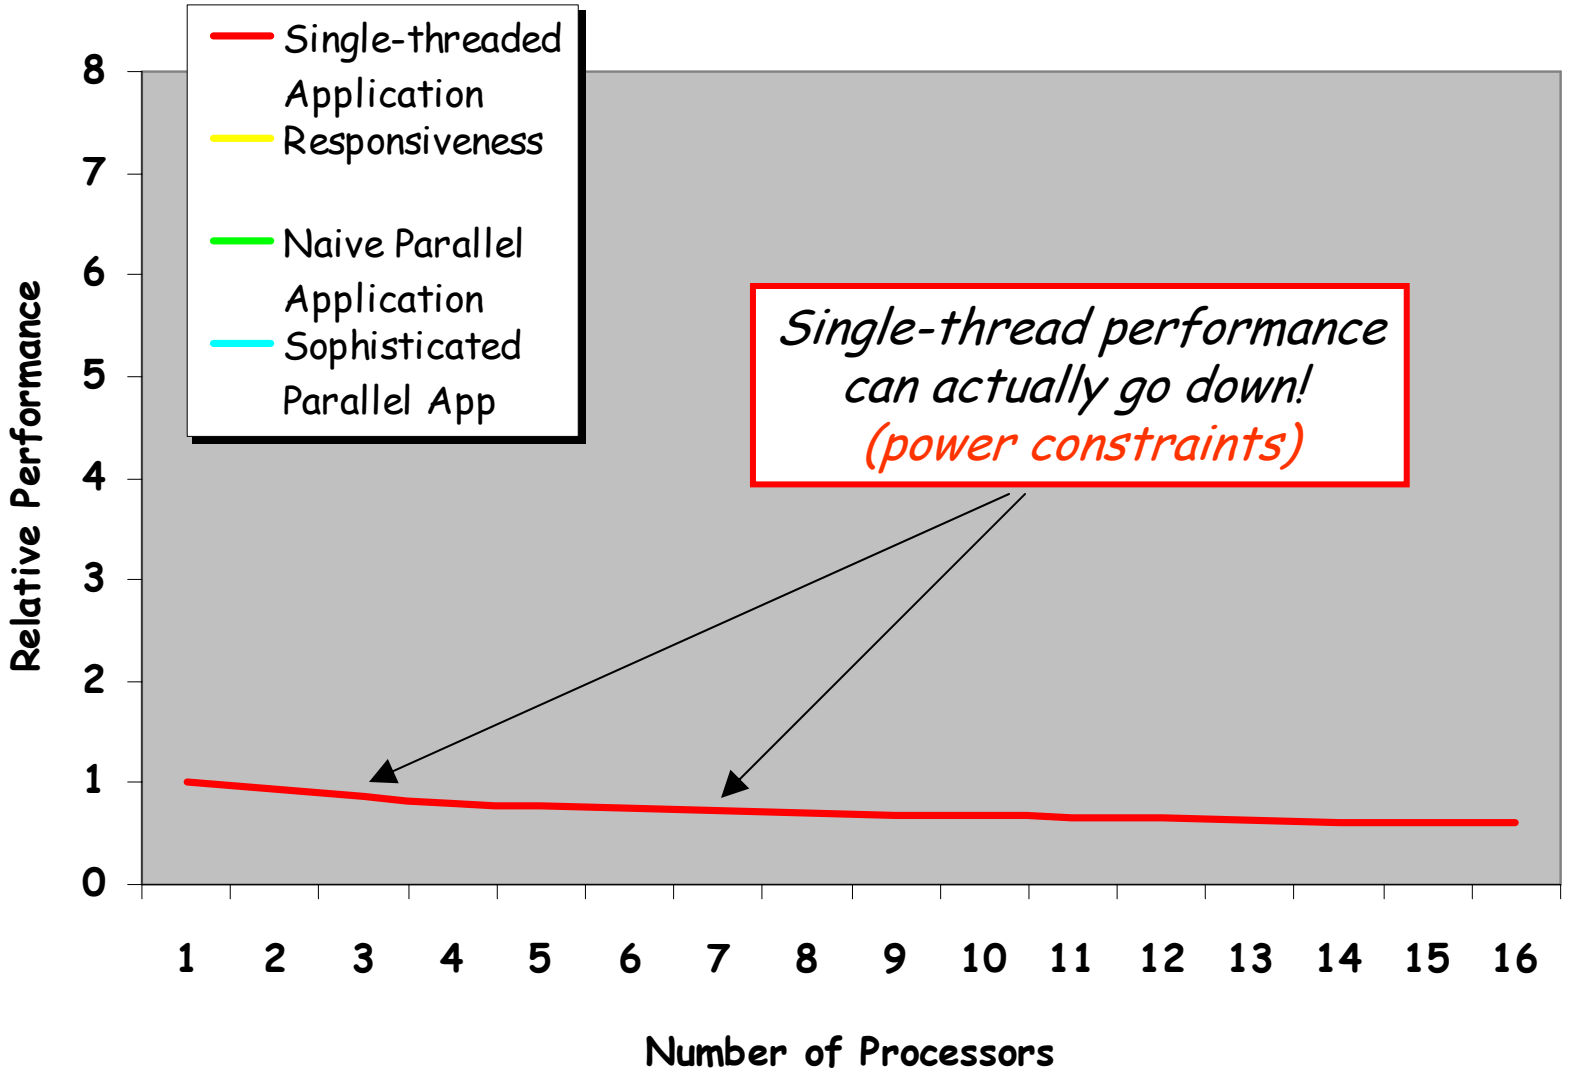

CMP Performance *(Theoretical values)*

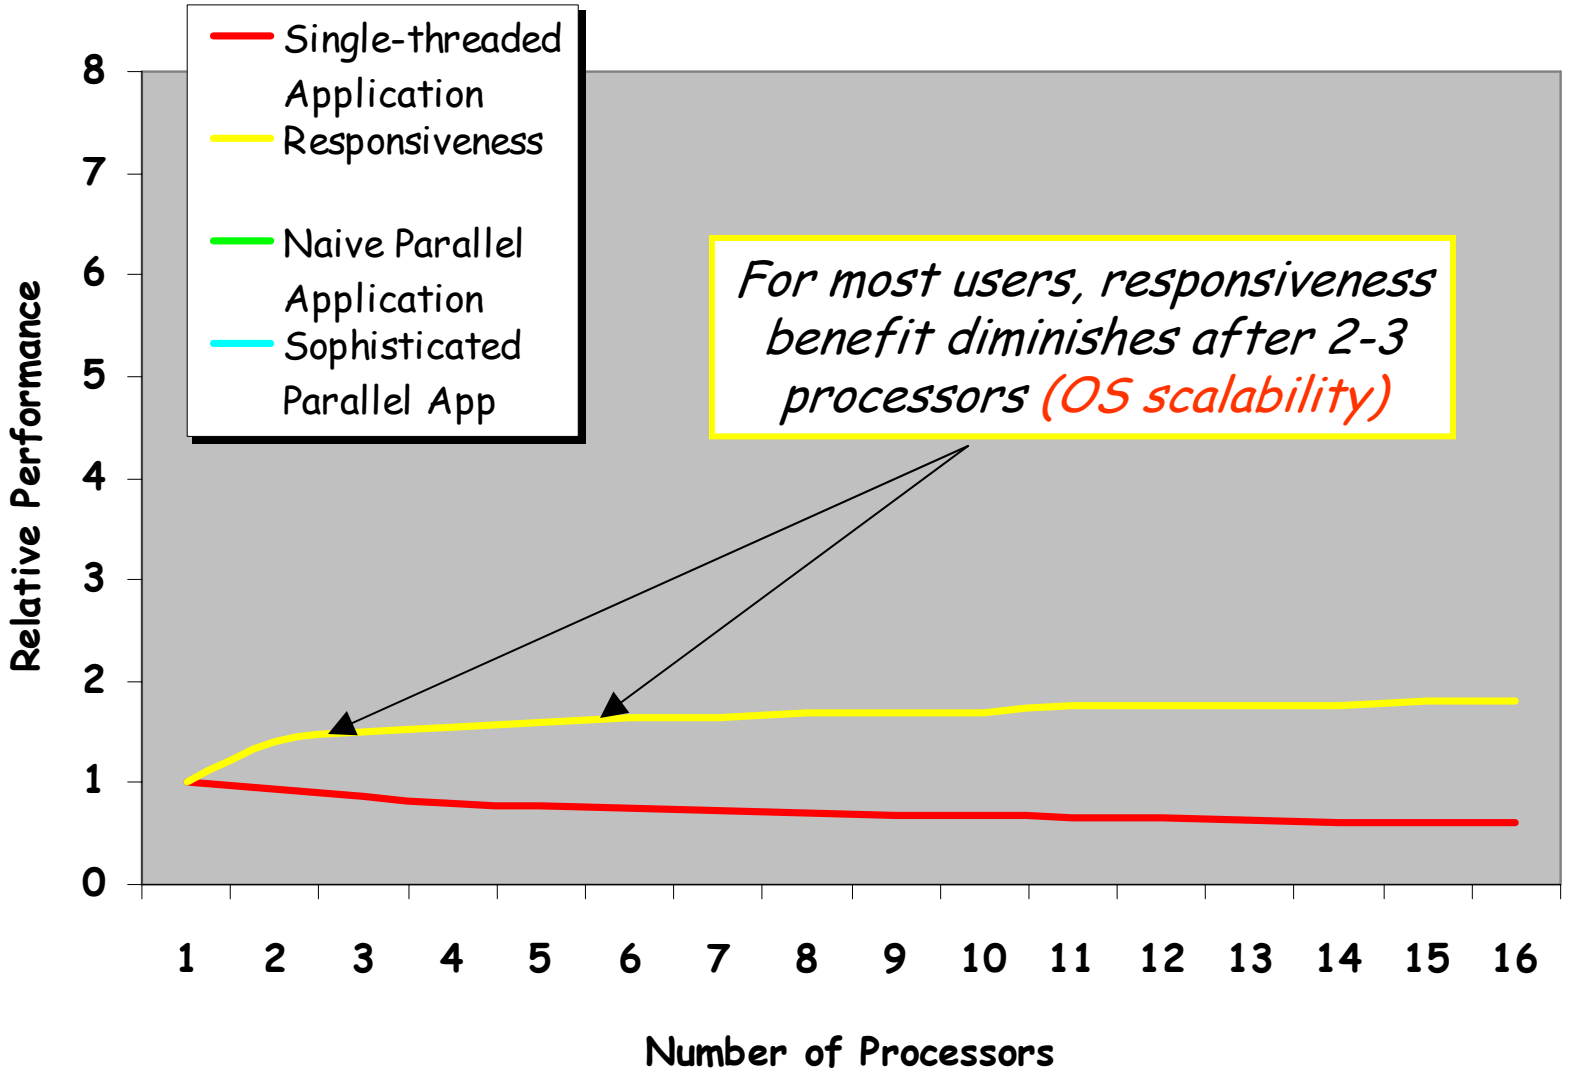

CMP Performance *(Theoretical values)*

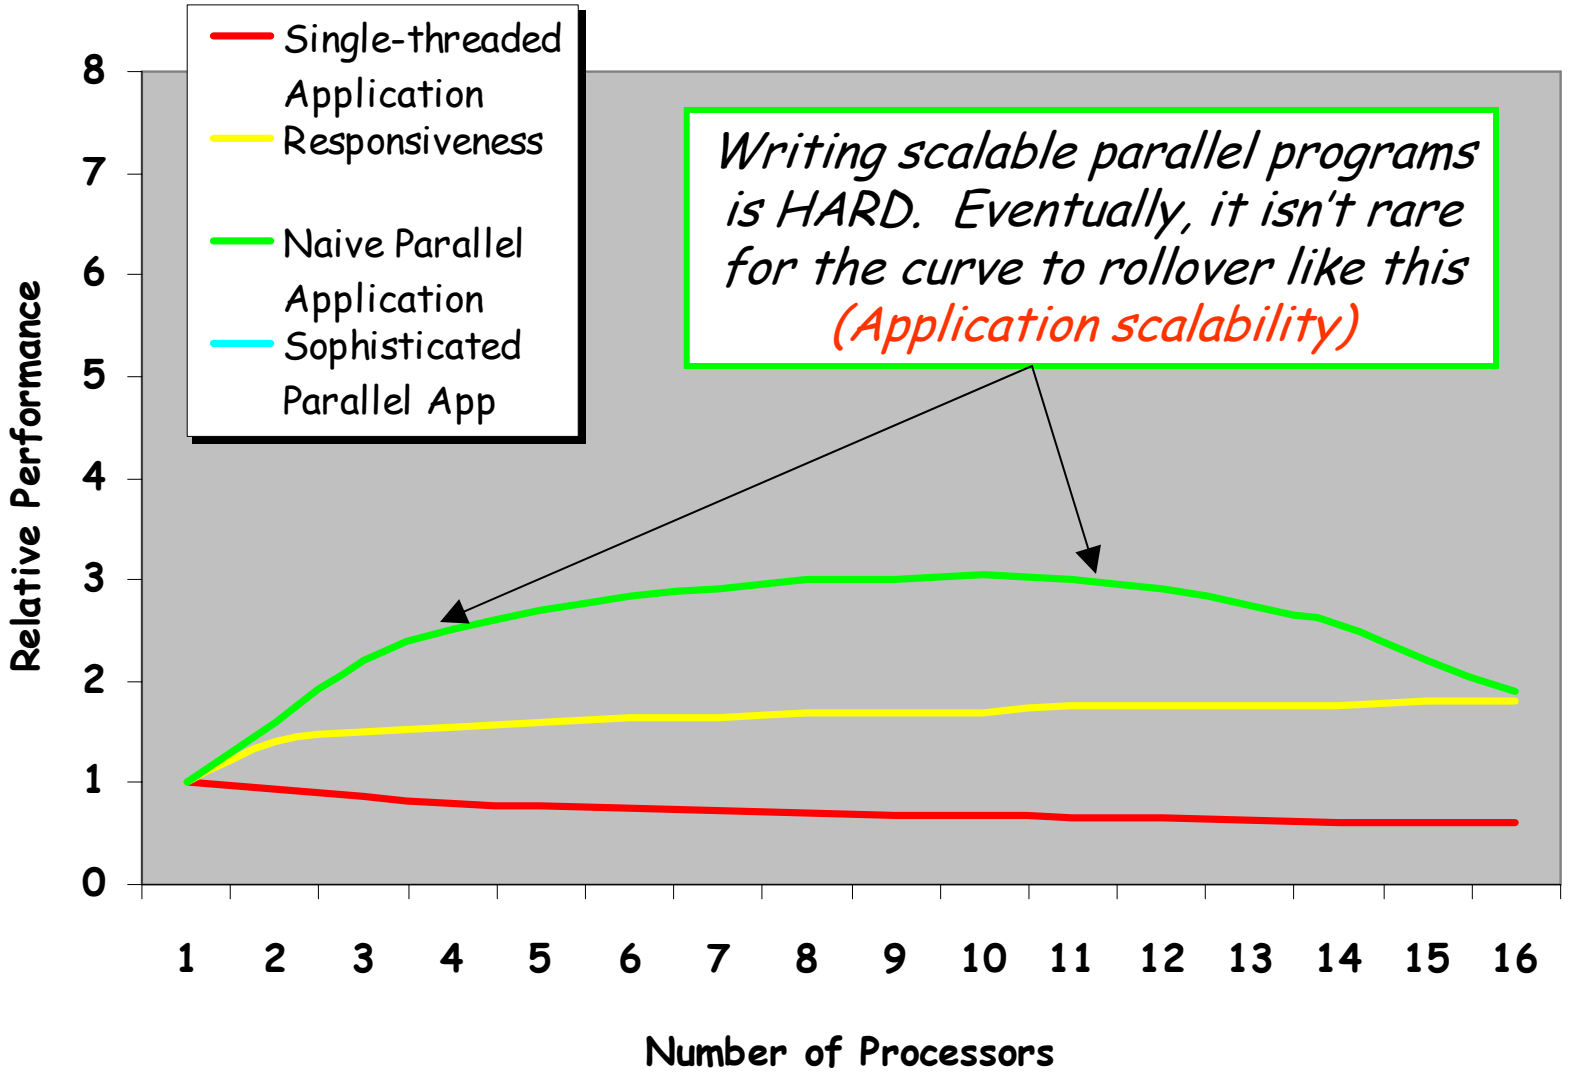

CMP Performance *(Theoretical values)*

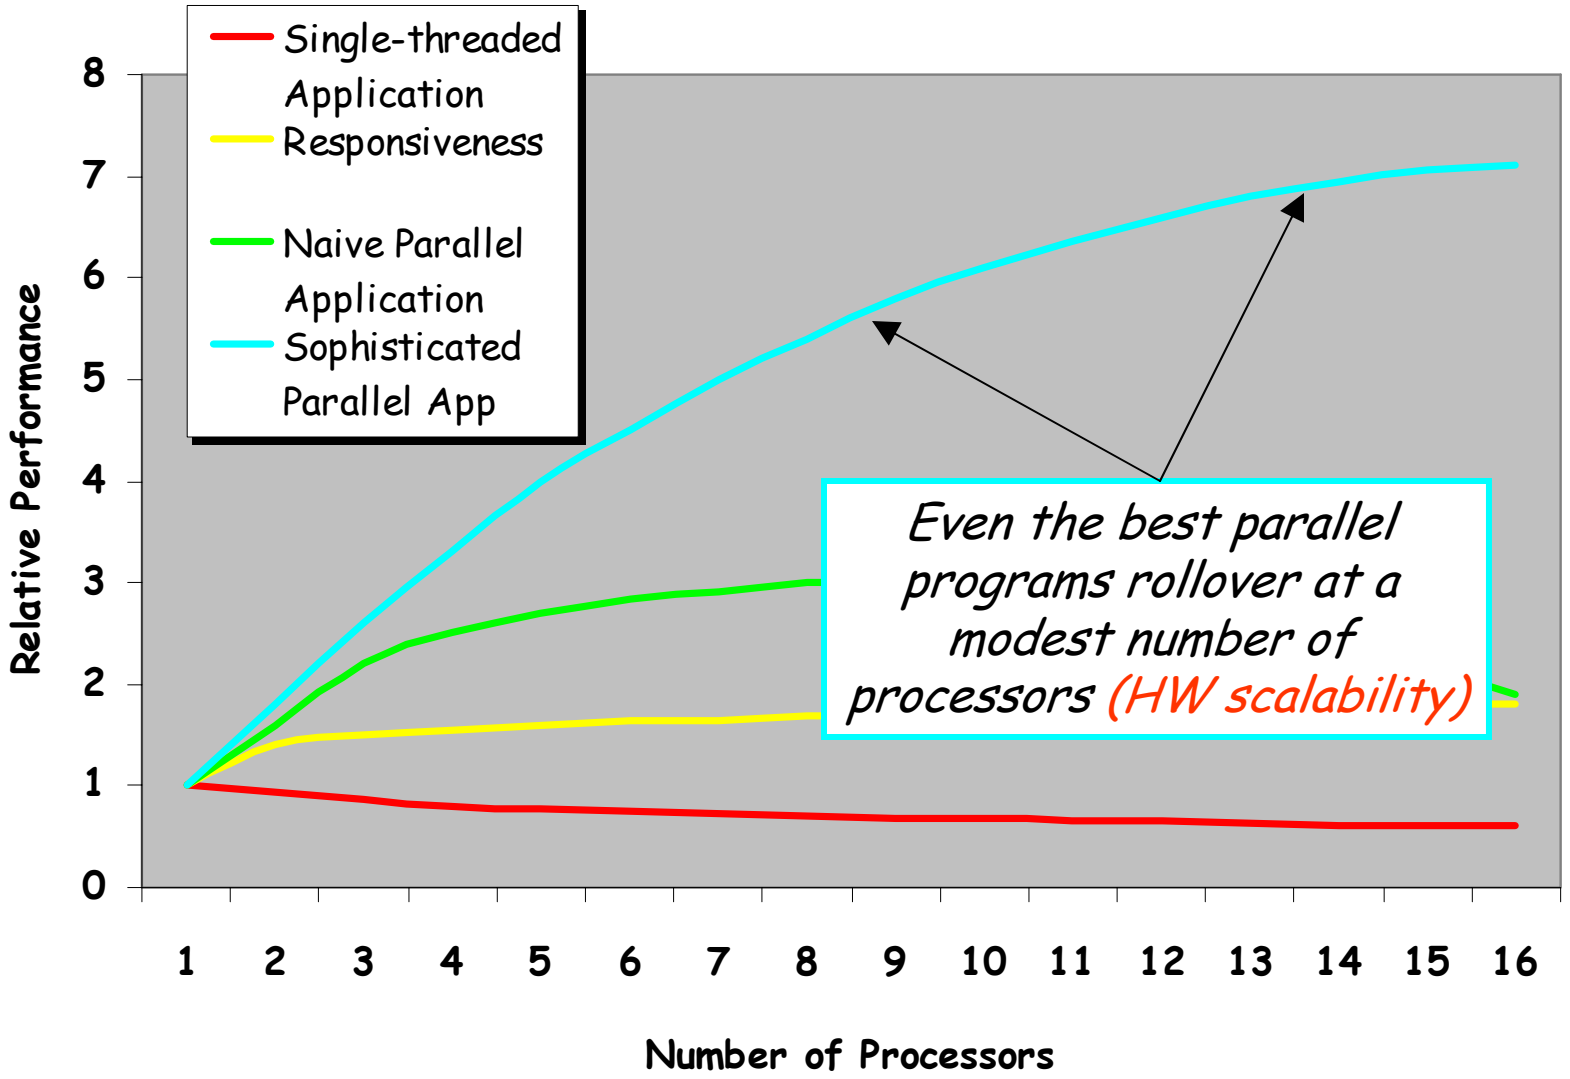

The Transition to Parallel Applications

Single-threaded Applications

- Most of today's applications
- Well understood optimization techniques
- Advanced development, analysis and debug tools
- Conceptually, easy to think about

Parallel Applications

- Small number of applications (worked by experts for 10+ yrs)
- Awkward development, analysis and debug environments
- Parallel programming is hard!
- Amdahl's law is still a law…
- SW productivity is already in a crisis → ***this worsens things!***

Establishing an *appropriate balance* is key for managing this important transition

Understanding Compatibility and Transitions

- The compatibility requirement takes many forms:
  - ISA compatibility (x86)
  - Structural compatibility (single-thread programs)
  - Socket compatibility (upgradeable systems)

- Will this ever change?
  - Much slower than what some people want you to believe
  - x86 acceptance has eclipsed most RISC architectures
  - Our history shows:
    - *Viable products don't bet on wildly incompatible solutions*
    - *Actual results dominated by weakest link in the HW/SW symbiotic relationship*
- Eventually, higher-level abstractions will take over
  - Maturing of software architectures and *layers*
  - Waste some performance to enable more interoperability
  - More on this later in the talk …

Ideas unlikely to work

Sudden turns

The fleet approach

Forced turns

Reliability and Robustness

Q: Which system has better performance?

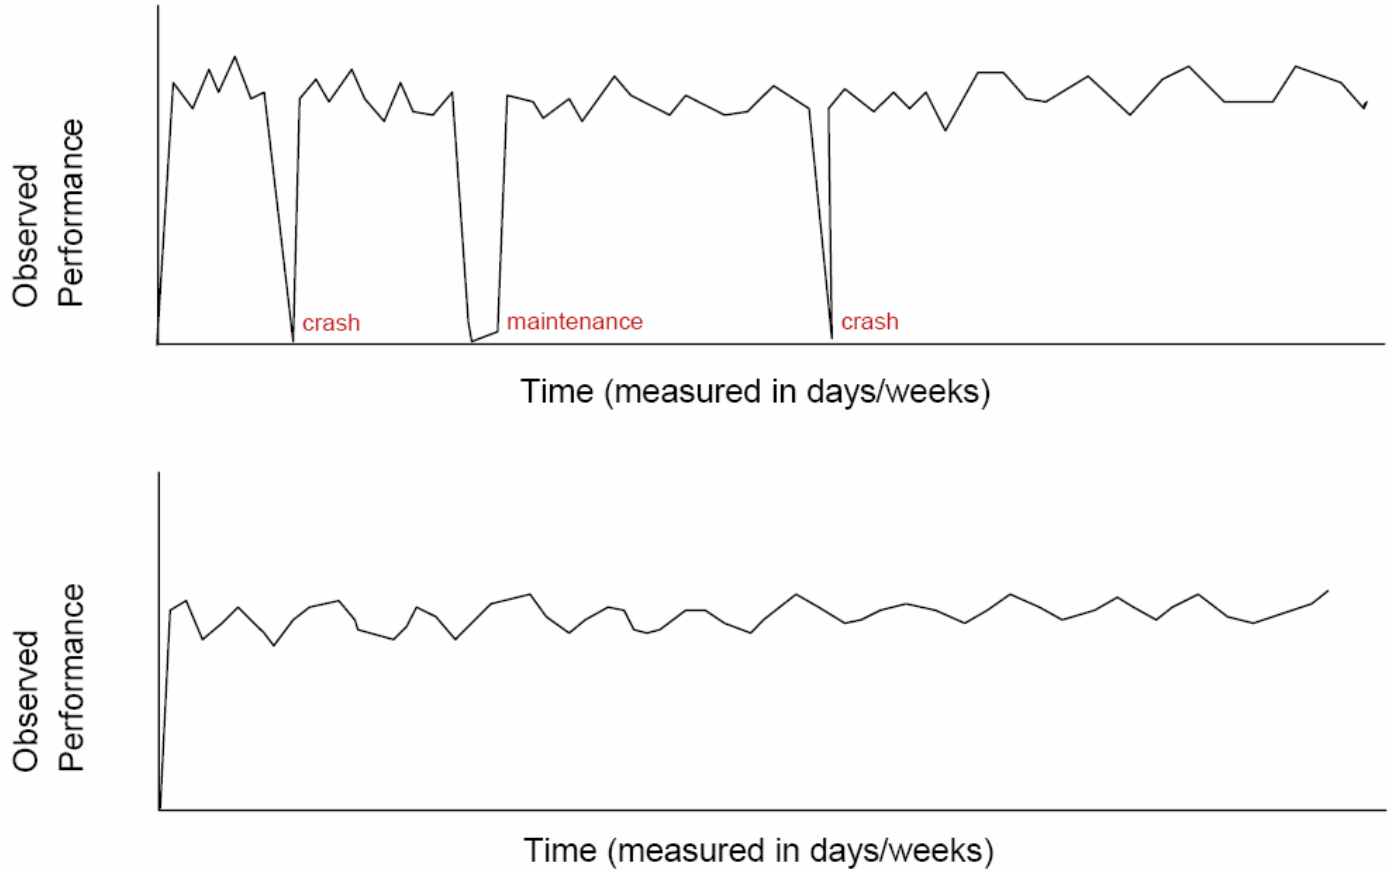

For many customers, this is more important than raw performance

Talk Outline

- What drives our industry?
- Elements of Balanced System Performance
  - System-level bandwidth and power consumption
  - Single thread versus multi-thread performance
  - Understanding compatibility and transitions
  - Reliability
- **AMD's Next Generation Processor Technology**
  - **Direct Connect Architecture enhancements**
  - **Core enhancements**
- Crystal Ball: *New trends that may shape the future*
- Summary

AMD's Next Generation Processor Technology

• Scalable performance and balance	Fast HyperTransport™ links (up to 5.2 GT/sec) Additional bandwidth enhancements On-chip shared L3 cache
• Maintain performance per watt leadership	Independent NB and CPU power management Independent CPU P-state and C-state controls
• Performance on diverse workloads	Enhanced IPC CPU core; >2X FPU performance 48-bit virtual and physical address space 1GB large page support Platform support for co-processors
• Compatibility	DDR2 memory support with migration to DDR3 FBDIMM Gen1 and Gen2 at the *appropriate time* HyperTransport™-1 backwards compatibility
• Enhanced Virtualization	I/O Virtualization Nested paging support
• Enhanced RAS	Memory mirroring Data poisoning support HyperTransport™ retry protocol support

AMD's Next Generation Processor Technology

Optimized for 65nm SOI and beyond

Native quad core die

Expandable shared L3 cache

IPC enhanced CPU cores

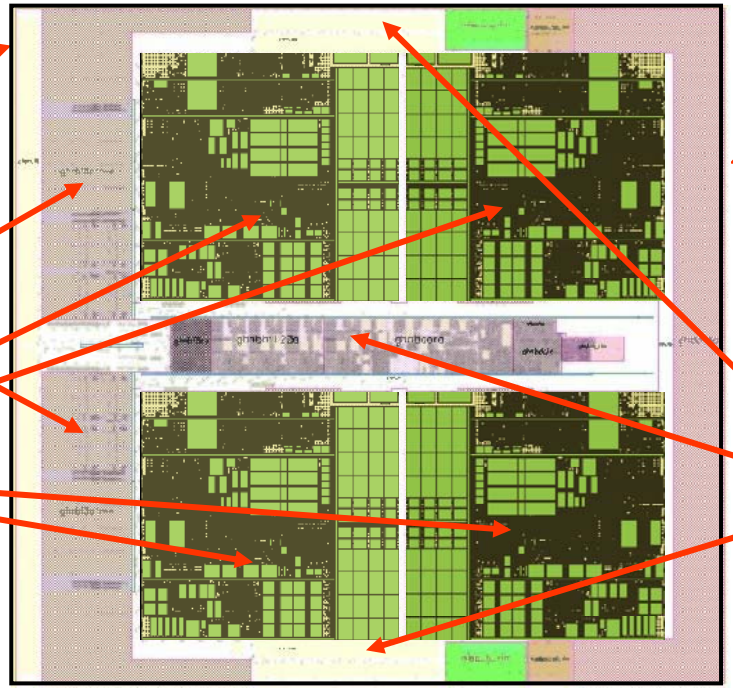

Enhanced Direct Connect Architecture and Northbridge

- 32B instruction fetch
- Improved branch prediction
- Out-of-order load execution
- Up to 4 DP FLOPS/cycle
- Dual 128-bit SSE dataflow
- Dual 128-bit loads per cycle
- Bit Manipulation extensions (LZCNT/POPCNT)
- SSE extensions (EXTRQ/INSERTQ, MOVNTSD/MOVNTSS)

- HyperTransport™ links (5.2GT/sec)
- Enhanced crossbar
- DDR2 with migration path to DDR3
- FBDIMM *when appropriate*
- Enhanced power management and RAS

Talk Outline

- What drives our industry?
- Elements of Balanced System Performance
  - System-level bandwidth and power consumption
  - Single thread versus multi-thread performance
  - Understanding compatibility and transitions
  - Reliability
- AMD's Next Generation Processor Technology
  - Direct Connect Architecture enhancements
  - Core enhancements
- **Crystal Ball: *New trends that may shape the future***
- **Summary**

1. Co-processors and Accelerators

Promising Concept

- Excellent way to get power-efficient performance boosts
  - Special-purpose, tuned solutions for common functions
  - Drop to low-power states when not in use
  - Enabled by Modern API's
- Aligns with modularity imperative
  - Co-processor becomes another (optional) "IP block"
  - Micro-architecture: *Command delivery, Synchronization, Streaming*
- Many possible opportunities now, and/or in the future
  - Media processing
  - JVM/CLR runtime hosting
  - NIC integration (TOE, XML, SSL, etc)

HyperTransport™ HTX™ Interface Enables System-level Coprocessing Today

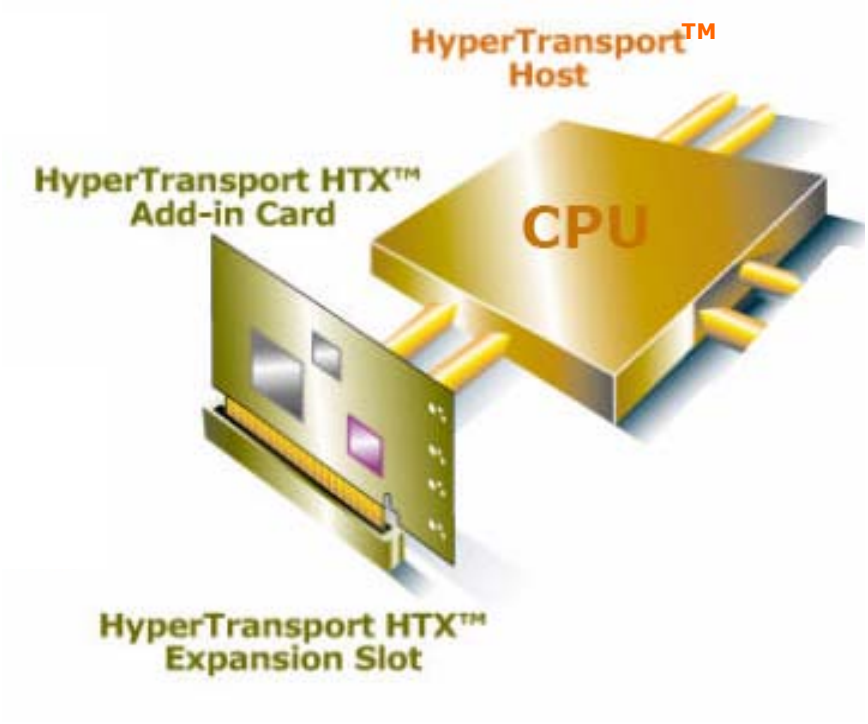

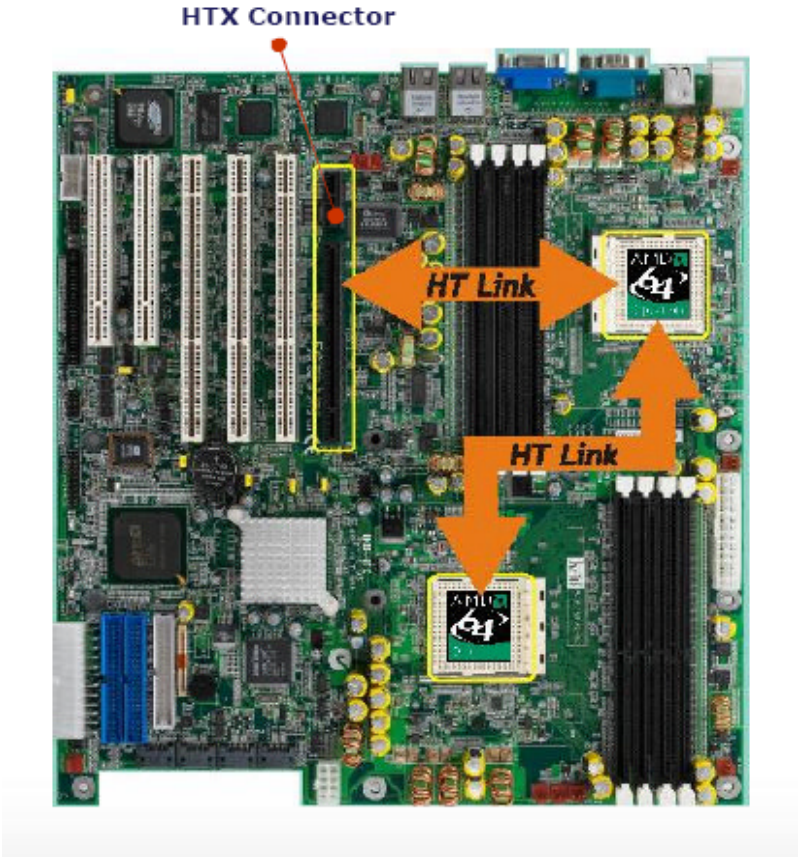

2. Chip Multiprocessor (CMP) Generations

- *Generation-0*
  - *Hasty integration of a bunch of cores*
  - *Jam multiple single core devices into the same package*
- ***Generation-1***
  - Design from the ***system-level inward***
  - Balanced integration of cores and system functionality
  - Address ***HW scalability*** inhibitors
- ***Generation-2***
  - Understand and address ***interference*** issues *(destructive & constructive)*
  - ***Chip-level framework*** & infrastructure to support range of CMP variants
  - Improve ***SW visibility*** & ***optimization*** for parallel application development
- ***Generation-3***
  - Introspective self-management
  - Runtime adaptive flexibility
    - *Specialized resources and optimizations*
    - *Engage resources appropriate for each application*

3. The role of abstraction the next phase of computing

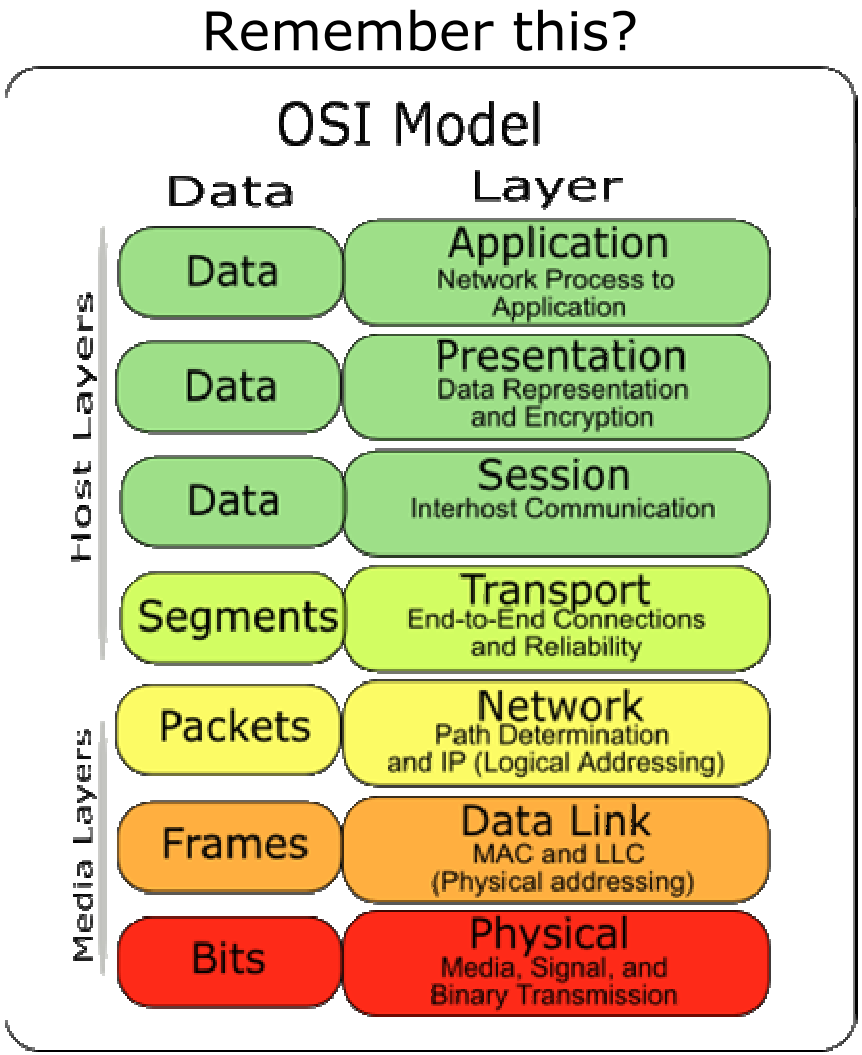

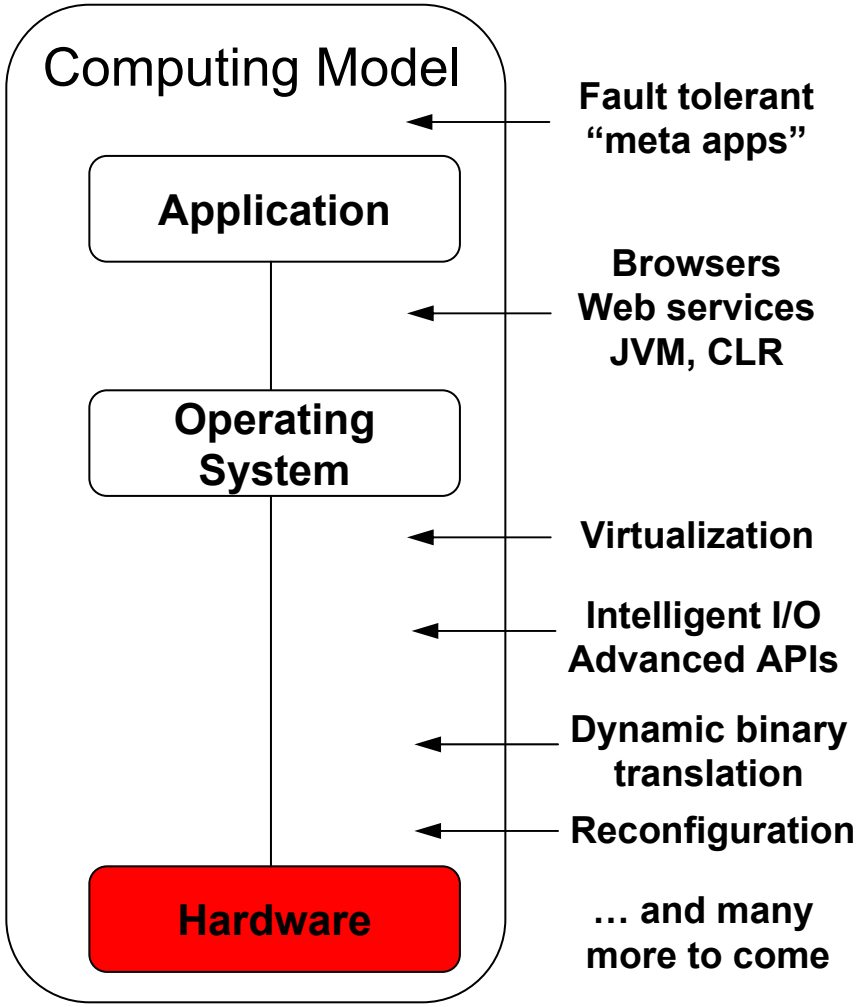

Foreshadows a HUGE paradigm shift !

AMD

Summary

- The semiconductor industry is driven by economics
  - Moore's Law is a key enabler in this cycle
  - But, underlying all of this is *customer value*
  - Today, customer value goes way beyond just performance
- AMD is focused on ***customer-centric innovation*** and ***value***
  - System-level balance and power efficiency
  - Appropriate balance of single-thread performance and throughput
- AMD is investing heavily in extending our leadership
  - Next generation CPU technology
  - Next generation Direct Connect Architecture technology
  - Developing a fundamental understanding of important emerging trends

Enjoy the conference !

www.amd.com/power